

Claims

What is claimed is:

1. A method for dynamic striping, comprising:
 - receiving a request to write a data block into a storage pool;
 - determining a physical disk location in the storage pool to store the data block using a dynamic striping policy;
 - storing the data block at the physical disk location; and
 - storing a first indirect block in the storage pool, wherein the first indirect block comprises the data block location and the data block checksum.
2. The method of claim 1, further comprising:
 - retrieving the data block using the first indirect block.
3. The method of claim 1, further comprising:
 - assembling the first indirect block, wherein assembling the first indirect block comprises populating a block pointer.
4. The method of claim 3, wherein populating the block pointer comprises:
 - storing the data block checksum in a checksum field within the block pointer; and
 - storing the data block location in the block pointer, wherein storing the data block location comprises storing a metaslab ID and offset.
5. The method of claim 4, further comprising:
 - storing a birth value in a birth field within the block pointer.
6. The method of claim 3, wherein the first indirect block is assembled using a data management unit.

7. The method of claim 1, wherein the storage pool comprises at least one storage device.
8. The method of claim 1, wherein the storage pool is divided into a plurality of metaslabs.
9. The method of claim 8, wherein each of the plurality of metaslabs is associated with a metaslab ID.
10. The method of claim 9, wherein the data block location comprises the metaslab ID and an offset.
11. The method of claim 1, wherein storing the data block comprises using a storage pool allocator.
12. The method of claim 1, wherein the dynamic striping policy comprises at least one selected from the group consisting of a dynamic striping policy based on physical disk speed, a dynamic striping policy based on free space available on physical disks, a dynamic striping policy based on load on physical disks, and a round robin policy.
13. A system for storing a data block, comprising:
 - a storage pool comprising the data block and a first indirect block, wherein the first indirect block comprises a data block checksum and a data block location; and
 - a storage pool allocator configured to store the data block and the first indirect block in the storage pool using a dynamic striping policy.
14. The system of claim 13, further comprising:
 - a second indirect block, comprising a first indirect block checksum and a first indirect block location,

wherein the storage pool allocator is further configured to store the second indirect block in the storage pool.

15. The system of claim 13, further comprising:

a data management unit configured to assemble the first indirect block and request the storage pool allocator store the first indirect block.

16. The system of claim 13, wherein the dynamic striping policy comprises at least one selected from the group consisting of a dynamic striping policy based on physical disk speed, a dynamic striping policy based on free space available on physical disks, a dynamic striping policy based on load on physical disks, and a round robin policy.

17. The system of claim 13, wherein the storage pool comprises at least one storage device.

18. The system of claim 13, wherein the storage pool is divided into a plurality of metaslabs.

19. The system of claim 18, wherein each of the plurality of metaslabs is associated with a metaslab ID.

20. The system of claim 19, wherein the data block location comprises the metaslab ID and an offset.

21. A computer system for dynamic striping, comprising: ↘

a processor;

a memory;

a storage device; and

software instructions stored in the memory for enabling the computer system under control of the processor, to:

receive a request to write a data block into a storage pool;
determine a physical disk location in the storage pool to store the data block using
a dynamic striping policy;
store the data block at the physical disk location; and
store a first indirect block in the storage pool, wherein the first indirect block
comprises the data block location and the data block checksum.

22. A network system having a plurality of nodes, comprising:

a storage pool comprising the data block and a first indirect block, wherein the
first indirect block comprises a data block checksum and a data block
location; and
a storage pool allocator configured to store the data block and the first indirect
block in the storage pool using a dynamic striping policy,
wherein the storage pool is located on any one of the plurality of nodes, and
wherein the storage pool allocator is located on any one of the plurality of nodes.